UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P O Box 1450
Alexandria, Virginia 22313-1450
www uspto gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/629,459	07/29/2003	Guogen Zhang	SVL920030024US1	7326

45729 7590 04/01/2009
GATES & COOPER LLP
6701 CENTER DRIVE WEST
SUITE 1050
LOS ANGELES, CA 90045

EXAMINER	
ADAMS, CHARLES D	

ART UNIT	PAPER NUMBER
2164	

MAIL DATE	DELIVERY MODE
04/01/2009	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

UNITED STATES PATENT AND TRADEMARK OFFICE

Commissioner for Patents
United States Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450
www.uspto.gov

BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

Application Number: 10/629,459
Filing Date: July 29, 2003
Appellant(s): ZHANG ET AL.

George H. Gates
For Appellant

EXAMINER'S ANSWER

This is in response to the appeal brief filed 18 December 2008 appealing from the Office action mailed 18 July 2008.

(1) Real Party in Interest

The real party in interest is IBM Corporation.

(2) Related Appeals and Interferences

The examiner is not aware of any related appeals, interferences, or judicial proceedings which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

(3) Status of Claims

The statement of the status of claims contained in the brief is correct.

(4) Status of Amendments After Final

The appellant's statement of the status of amendments after final rejection contained in the brief is correct.

(5) Summary of Claimed Subject Matter

The summary of claimed subject matter contained in the brief is correct.

(6) Grounds of Rejection to be Reviewed on Appeal

The appellant's statement of the grounds of rejection to be reviewed on appeal is correct.

(7) Claims Appendix

The copy of the appealed claims contained in the Appendix to the brief is correct.

(8) Evidence Relied Upon

"Implementation of Two Semantic Query Optimization Techniques in a DB2 Universal Database"	**Cheng et al.**	**9-1999**
US 5,963,936 A	**Cochrane et al.**	**10-1999**
US 6,438,741 B1	**Al-omari et al.**	**8-2002**

(9) Grounds of Rejection

The following ground(s) of rejection are applicable to the appealed claims:

Claim Rejections - 35 USC § 103

1. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

2. Claims 1-3 are rejected under 35 U.S.C. 103(a) as being unpatentable over Cheng et al. ("Implementation of Two Semantic Query Optimization Techniques in DB2

Universal Database"), in view of Cochrane et al. (US Patent 5,963,936), and further in view of Al-omari et al. (US Patent 6,438,741).

As to claim 1, Cheng et al. teaches a method of optimizing a query in a computer system, the query being performed by the computer system to retrieve data from a database stored on the computer system (see Abstract), the method comprising:

(a) during compilation of the query, maintaining a GROUP BY clause (see Cheng et al. Page 1, Example 1, and Page 5, query 1)

Cheng et al. does not teach with one or more GROUPING SETS, ROLLUP or CUBE operations

Cocharane et al. teaches with one or more GROUPING SETS, ROLLUP or CUBE operations (see column 7, lines 26-30, and column 7, lines 44-48)

Cheng et al. as modified teaches in its original form, instead of rewriting the GROUP BY clause, until after query rewrite (see Cheng et al. Page 1, Example 1, and Page 5, query 1. In Q'1, the group by clause has been retained); and

(b) at a later stage of query compilation, translating the GROUP BY clause with the GROUPING SETS, ROLLUP, or CUBE operations into a plurality of levels, wherein each of the levels has one or more grouping sets (see Cochrane et al. 8:26-42, Figure 7. This step occurs after the step listed above) comprised of grouping columns (see 11:62-12:15. The GROUP BY sets are comprised of columns a, b, x, and y),

Cheng et al. as modified does not teach generating a query execution plan for the query with a super group block having an array of pointers, wherein each pointer points to the grouping sets for a particular one of the levels.

Al-omari et al. teaches generating a query execution plan for the query with a super group block having an array of pointers, wherein each pointer points to the grouping sets for a particular one of the levels (see 10:36-48, 14:28-35, 41-43. Also see Figure 15B, 9:31-43 and 16:50-59);

Cheng et al. as modified teaches:

(c) performing the query execution plan to retrieve data from a database stored on the computer system (see Cochrane et al. 7:41-43).

Therefore, it would have been obvious to one of ordinary skill in the art at the time the invention was made to have modified Cheng et al. by the teachings of Cochrane et al., since Cochrane et al. teaches that "a method for detecting and stacking grouping sets to support group by operations with grouping sets, rollup, and cube extensions in relational database management systems, with greatly reduced numbers of grouping sets" (see Abstract).

Therefore, it would have been obvious to one of ordinary skill in the art at the time the invention was made to have further modified Cheng et al. by the teachings of Al-omari et al., since Al-omari et al. teaches "a system and method for optimizing complex SQL database queries" (see 3:18-19).

As to claim 2, Cheng et al. as modified teaches further comprising:

(1) at query execution time, dynamically determining a grouping sequence for the GROUP BY clause with the GROUPING SETS, ROLLUP or CUBE operations based on intermediate grouping sets, in order to optimize the grouping sets sequence (see Cochrane et al. 8:26-42, Figure 7).

As to claim 3, Cheng et al. as modified teaches wherein the dynamically determining step further comprises (1) performing a GROUP BY for a base grouping set and then optimizing execution of the grouping sets sequence by selecting a grouping set having lowest cardinality from a previous one of the levels as an input to a grouping set on a next one of the levels (see Cochrane et al. 11:43-47. The GROUP BYs “are stacked from greatest to least cardinality”. There is only one grouping set per level. It is inherent, then, that the chosen grouping set sequence from a previous one of the levels will be the smallest one on that grouping set’s level), and (2) performing a UNION ALL operation on the grouping sets (see Cochrane et al. 11:47-49 and Figure 7. “The base group by and all the GROUP BYs for ROLLUP1 are unioned together”. If all of the GROUP BYs are unioned together, then it is functionally equivalent to a UNION ALL).

(10) Response to Argument

As to the independent claim, Appellant argues that “the combination of Cheng, Cochrane, and Al-omari does not teach or suggest the limitations of the independent claim, and goes on to state that "the query of Cheng merely illustrates the technique of

join elimination in the context of semantic query optimization. However, Cheng merely shows the GROUP BY clause in the same form in both the original query and the optimized query, indicating that the GROUP BY clause is not "maintained during compilation" and then "translated at a later stage of query compilation," as recited in Appellants' claims. Instead, the GROUP BY clause of Cheng is apparently left untouched during the join elimination optimization. Moreover, the GROUP BY clause in Cheng does not have GROUPING SETS, ROLLUP, or CUBE operations in its original form, as admitted by the Office Action. Therefore, the query of Cheng, and the example cited by the Office Action, have no relevance to Appellants' claims".

In response to this argument, it is noted that Cheng et al. teaches wherein two stages of query optimization exist, “query rewrite optimization”, and “query plan optimization” (see page 3, section 2, 1st paragraph), and that “after an input query is parsed and converted to an intermediate form called query graph model, the graph is transformed by the Query Rewrite Engine into a logically equivalent but more efficient form using heuristics” (see page 3, section 2, second paragraph). Cheng et al. also states that "The design of the query rewrite engine is ideal for the implementation of SQO" (see page 3, column 2, 1st paragraph after Example 3).

A GROUP BY clause is included in "Query 1" of Cheng et al. (see page 2, Example 1). This “Query 1” is then rewritten using the SQO techniques, to result in Q'_1 (see page 5, “Query 1”). “Query 1” contains a “group by” clause. This “group by” clause is maintained while the query is being processed by the “query rewrite engine” using the SQO techniques of Cheng et al., and still exists unchanged in rewritten query Q'_1.

Therefore, the “group by” clause is maintained in its original form, instead of being rewritten, until after query rewrite.

Examiner relied upon Cochrane et al. to teach wherein the query contained one or more GROUPING SETS, ROLLUP, or CUBE operations (see Cochrane et al. 7:26-30, and 7:44-48).

Appellant argues in regards to Cochrane et al. that “This optimization scheme of Cochrane says nothing about maintaining a GROUP BY clause with one or more GROUPING SETS, ROLLUP, or CUBE operations in its original form, instead of rewriting the GROUP BY clause, until after query rewrite. Instead, the optimization scheme of Cochrane reduces the GROUP BYs during query rewrite, which means that the GROUP BY clause is not maintained in its original form until after query rewrite, but instead the GROUP BY clause is rewritten”. In response to this argument, it is first noted that the rejection was an obviousness rejection in view of Cheng et al., in view of Cochrane et al., and in view of Al-Omari et al., and was based on 35 U.S.C. 103(a).

Cheng et al. first submits a QGM to a “query rewrite engine”. Cheng et al. also mentions “query rewrite optimization” and “query plan optimization” (see page 3, column 1). Cochrane et al. receives a query, forms a "query graph model", and rewrites the query, and submit this QGM to a query plan optimizer (see 7:22-35). It is noted that the GROUP BYs in Cochrane et al. are maintained in their original form until the moment of query rewrite. They no longer exist in their original form after query rewrite. Thus,

Cochrane et al. does teach "maintaining a GROUP BY clause in its original form until after query rewrite".

In regards to Al-omari et al., Appellants argue that “the groups in Al-omari are in no way equivalent to Appellants’ claimed super group block. Specifically, the memo structure of Al-omari includes one or more groups, where each group contains an array of pointers to one or more logical expressions, an array of pointers to one or more physical expressions, an array of pointers to one or more contexts, an array of pointers to one or more plans, and an exploration pass indicator. In Appellants’ claims, on the other hand, the super group block supports the translation of a GROUP BY clause with the GROUPING SETS, ROLLUP or CUBE operations into a plurality of levels, wherein each of the levels has one or more grouping sets comprised of grouping columns, the super group block has an array of pointers, and each pointer of the super group block points to the grouping sets for a particular one of the levels. This super group block of Appellants’ claims recites different structure and functions as compared to the memo structure of Al-omari”.

In response to this argument, it is noted that the memo structure of Al-omari et al. is a super group block (see 10:36-48, 14:28-35 and 41-43). The “groups” of the memo structure contain pointers to logical and physical expressions. It is also noted that each group of the memo structure contain pointers to a table expression (see Figure 15B, 9:31-34, and 16:50-59. Table expressions include “group bys”. While only “joins” and “scans” are listed in Figure 15B, 9:31-43 and 16:50-59 indicate that “group bys” are also

valid table expressions). It is also noted that the Table expressions pointed to by the groups may accept as input other groups (see 24:38-54). Cochrane et al. teaches to generate a stack of GROUP BYs, in which each previous GROUP BY serves as an input to the next GROUP BY (see 11:34-12:35). As Al-omari et al. teaches a super group block Memo structure with pointers to various table expressions (wherein "GROUP BY" is a valid table expression), wherein each table expression may serve as input to another table expression, and as Cochrane et al. teaches creating a stack of different levels of GROUP BYs, wherein results from GROUP BYs on one level are input into the GROUP BYs on the next, it would have been obvious to one of ordinary skill in the art to have combined the two to have made the currently claimed invention.

In regards to claim 2, Appellant argues that "Cochrane merely describes constructing a query graph model (QGM) that includes a stacking of GROUP BYs. However, Cochrane says nothing about group sets sequences, or dynamically determining, at query execution time, a grouping sets sequence for a GROUP BY clause with GROUPING SETS, ROLLUP, or CUBE operations, based on intermediate grouping sets, in order to optimize the grouping sets sequence." In response to this argument, it is noted that the grouping sets of Cochrane et al. are stacked (see 8:26-42). This is a sequence of grouping sets. It is also noted that this is determined at runtime, after the query is parsed (see 7:22-34). Thus, it is dynamic. Finally, the stack of GROUP BYs of Cochrane et al. are setup such that there is a base GROUP BY which

serves as an input to the GROUP BY above it. This step is done iteratively (see 11:34-12:35).

In regards to claim 3, Appellant argues that "Cochrane merely describes the base step in the construction of a stack of GROUP BYs for a ROLLUP, wherein a base GROUP BY and all other GROUP BYs for the ROLLUP are "unioned" together. However, Cochrane says nothing about grouping sets sequences, or performing a GROUP BY for a base grouping set and then optimizing execution of the grouping sets sequence by selecting a grouping set having lowest cardinality from a previous one of the levels as an input to a grouping set on a next one of the levels, and then performing a UNION ALL operation on the grouping sets". As stated above, the GROUP BYs of Cochrane et al. are stacked (see 8:26-42). This is a sequence of grouping sets. A GROUP BY is performed on a base grouping set (see 11:35-49). The GROUP BYs "are stacked from greatest to least cardinality". There is only one grouping set per level. It is inherent, then, that the chosen grouping set sequence from a previous one of the levels will be the smallest one on that grouping set's level. It is also noted that Cochrane et al. performs a UNION on all of the grouping sets as a final operation (see 11:35-61)

(11) Related Proceeding(s) Appendix

No decision rendered by a court or the Board is identified by the examiner in the Related Appeals and Interferences section of this examiner's answer.

For the above reasons, it is believed that the rejections should be sustained.

Respectfully submitted,

/Charles D Adams/

Charles D. Adams

Examiner

Art Unit 2164

Conferees:

/Charles Rones/

Supervisory Patent Examiner, Art Unit 2164

/Pierre M. Vital/
Supervisory Patent Examiner, Art Unit 2169

March 27, 2009